EAST Search History

Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
L1	151	generate near (bit adj stream) and "7"??.clas.	US-PGPUB; USPAT	OR	OFF	2007/09/10 10:57
L2	0	generate near (bit adj stream) and "7"??.clas. and tree and (child adj node)	US-PGPUB; USPAT	OR	OFF	2007/09/10 10:57
L3	0	generate near (bit adj stream) and "7"??.clas. and tree and (parent adj node)	US-PGPUB; USPAT	OR	OFF	2007/09/10 10:57
L4	18	generate near (bit adj stream) and "7"??.clas. and tree	US-PGPUB; USPAT	OR	OFF	2007/09/10 10:59
L5	3	generate near (bit adj stream) and "707".clas. and tree	US-PGPUB; USPAT	OR	OFF	2007/09/10 11:01
L6	12	(bit adj stream) and "707".clas. and tree and (parent adj node)	US-PGPUB; USPAT	OR	OFF	2007/09/10 11:00
L7	0	generate near (bit adj stream) and "709".clas. and tree	US-PGPUB; USPAT	OR	OFF	2007/09/10 11:01
L8	18	generate near (bit adj stream) and "7"??.clas. and tree	US-PGPUB; USPAT	OR	OFF	2007/09/10 11:01
L9	3	generate near (bit adj stream) and "7"??.clas. and tree and node	US-PGPUB; USPAT	OR	OFF	2007/09/10 11:01
L10	5397	707/102.ccls.	US-PGPUB; USPAT	OR	OFF	2007/09/10 11:15
L11	262	707/102.ccls. and @ay>"2005"	US-PGPUB; USPAT	OR	OFF	2007/09/10 11:16
L12	8	707/102.ccls. and @ay>"2005" and (bit stream).clm.	US-PGPUB; USPAT	OR	OFF	2007/09/10 12:19
L13	0	707/102.ccls. and @ay>"2005" and (generate near (bit adj stream)).clm.	US-PGPUB; USPAT	OR	OFF	2007/09/10 11:16
L14	0	707/102.ccls. and @ay>"2005" and (bit adj stream).clm.	US-PGPUB; USPAT	OR	OFF	2007/09/10 11:16
L15	5	(US-20060253448-$ or US-20030177341-$ or US-20040133848-$ or US-20070156750-$).did. or (US-5153591-$).did.	US-PGPUB; USPAT	OR	OFF	2007/09/10 11:18
L16	69925	heuer$.in. nad "707".clas.	US-PGPUB; USPAT	OR	OFF	2007/09/10 12:14
L17	15	heuer$.in. and "707".clas.	US-PGPUB; USPAT	OR	OFF	2007/09/10 12:15
L18	4	heuer$.in. and "707".clas. and (bit adj stream).clm.	US-PGPUB; USPAT	OR	OFF	2007/09/10 12:15

EAST Search History

L19	27	707/100.ccls. and @ay>"2005" and (bit stream).clm.	US-PGPUB; USPAT	OR	OFF	2007/09/10 12:20
L20	0	707/100.ccls. and @ay>"2005" and (bit adj stream).clm.	US-PGPUB; USPAT	OR	OFF	2007/09/10 12:20

Terms used: **index tree bit stream** Found **42,666** of **210,468**

Sort results by relevance Save results to a Binder Try an Advanced Search

Display results expanded form Search Tips Try this search in The ACM Guide

Open results in a new window

Results 1 - 20 of 200 Result page: **1** 2 3 4 5 6 7 8 9 10 next

Best 200 shown Relevance scale

1 The bits between the lambdas: binary data in a lazy functional language

Malcolm Wallace, Colin Runciman

October 1998 **ACM SIGPLAN Notices , Proceedings of the 1st international symposium on Memory management ISMM '98**, Volume 34 Issue 3

Publisher: ACM Press

Full text available: pdf(1.56 MB) Additional Information: full citation, abstract, references, citings, index terms

For the programmer, storage media are usually assumed to have a minimum atomic unit of transfer of one byte. However, sometimes it is useful to have an even finer storage granularity of one bit, for instance in order to compress data.This paper describes an API in the lazy functional language Haskell for treating storage media as arbitrary-length streams of bits without byte-alignment constraints. So far as possible, storage media are treated uniformly. In particular, bit-stream memory and binar ...

2 Configuration compression for FPGA-based embedded systems

Andreas Dandalis, Viktor K. Prasanna

February 2001 **Proceedings of the 2001 ACM/SIGDA ninth international symposium on Field programmable gate arrays FPGA '01**

Publisher: ACM Press

Full text available: pdf(203.25 KB) Additional Information: full citation, abstract, references, citings, index terms

FPGAs are a promising technology for developing high-performance embedded systems. The density and performance of FPGAs have drastically improved over the past few years. Consequently, the size of the configuration bit-streams has also increased considerably. As a result, the cost-effectiveness of FPGA-based embedded systems is significantly affected by the memory required for storing various FPGA configurations. This paper proposes a novel compression technique that reduces the memory requ ...

3 Course 4: State of the art in massive model visualization: Massive model visualization using realtime ray tracing

Dave Kasik

August 2007 **ACM SIGGRAPH 2007 courses SIGGRAPH '07**

Publisher: ACM Press

Full text available: pdf(34.96 MB) Additional Information: full citation, references

4 The feasibility of supporting large-scale live streaming applications with dynamic application end-points

Kunwadee Sripanidkulchai, Aditya Ganjam, Bruce Maggs, Hui Zhang

August 2004 **ACM SIGCOMM Computer Communication Review , Proceedings of the 2004 conference on Applications, technologies, architectures, and protocols for computer communications SIGCOMM '04**, Volume 34 Issue 4

Publisher: ACM Press

Full text available: pdf(461.96 KB) Additional Information: full citation, abstract, references, citings, index terms

While application end-point architectures have proven to be viable solutions for large-scale distributed applications such as distributed computing and file-sharing, there is little known about its feasibility for more bandwidth-demanding applications such as live streaming. Heterogeneity in bandwidth resources and dynamic group membership, inherent properties of application end-points, may adversely affect the construction of a usable and efficient overlay. At large scales, the problems become ...

Keywords: application-level multicast, live streaming, overlay multicast, peer-to-peer

5 When indexing equals compression: Experiments with compressing suffix arrays and applications

Luca Foschini, Roberto Grossi, Ankur Gupta, Jeffrey Scott Vitter

October 2006 **ACM Transactions on Algorithms (TALG)**, Volume 2 Issue 4

Publisher: ACM Press

Full text available: pdf(351.20 KB) Additional Information: full citation, abstract, references, index terms

We report on a new experimental analysis of high-order entropy-compressed suffix arrays, which retains the theoretical performance of previous work and represents an improvement in practice. Our experiments indicate that the resulting text index offers state-of-the-art compression. In particular, we require roughly 20% of the original text size---without requiring a separate instance of the text. We can additionally use a simple notion to encode and decode block-sorting transforms (such a ...

Keywords: Burrows--Wheeler Transform, Entropy, suffix array, text indexing

6 Progressive forest split compression

Gabriel Taubin, André Guéziec, William Horn, Francis Lazarus

July 1998 **Proceedings of the 25th annual conference on Computer graphics and interactive techniques SIGGRAPH '98**

Publisher: ACM Press

Full text available: pdf(2.53 MB) Additional Information: full citation, references, citings, index terms

Keywords: algorithms, geometric compression, graphics

7 Research session: database architectures for new hardware: Improving database performance on simultaneous multithreading processors

Jingren Zhou, John Cieslewicz, Kenneth A. Ross, Mihir Shah

August 2005 **Proceedings of the 31st international conference on Very large data bases VLDB '05**

Publisher: VLDB Endowment

Full text available: pdf(271.13 KB) Additional Information: full citation, abstract, references, citings, index

terms

Simultaneous multithreading (SMT) allows multiple threads to supply instructions to the instruction pipeline of a superscalar processor. Because threads share processor resources, an SMT system is inherently different from a multiprocessor system and, therefore, utilizing multiple threads on an SMT processor creates new challenges for database implementers.We investigate three thread-based techniques to exploit SMT architectures on memory-resident data. First, we consider running independent ope ...

8 Algorithm 806: SPRNG: a scalable library for pseudorandom number generation

Michael Mascagni, Ashok Srinivasan
September 2000 **ACM Transactions on Mathematical Software (TOMS)**, Volume 26 Issue 3
Publisher: ACM Press

Full text available: pdf(158.69 KB) Additional Information: full citation, appendices and supplements, abstract, references, cited by, index terms

In this article we present background, rationale, and a description of the Scalable Parallel Random Number Generators (SPRNG) library. We begin by presenting some methods for parallel pseudorandom number generation. We will focus on methods based on parameterization, meaning that we will not consider splitting methods such as the leap-frog or blocking methods. We describe, in detail, parameterized versions of the following pseudorandom number generators: (i) linear congruential generators, ...

Keywords: lagged-Fibonacci generator, linear congruential generator, parallel random-number generators, random-number software, random-number tests

9 Raster data, data structures and computational geometry: Regular and irregular multi-resolution terrain models: a comparison

Leila De Floriani, Paola Magillo
November 2002 **Proceedings of the 10th ACM international symposium on Advances in geographic information systems GIS '02**
Publisher: ACM Press

Full text available: pdf(240.42 KB) Additional Information: full citation, abstract, references, citings, index terms

The paper deals with the problem of modeling large-size terrain data sets. To this aim, we consider multi-resolution models based on triangle meshes. We analyze and compare two multi-resolution terrain models based on regular and irregular meshes. The two models are viewed as instances of a common multi-resolution model, that we call a *multi-resolution triangle mesh*. Our comparison takes into account the space requirements of the data structures implementing the two models as well their e ...

Keywords: multi-resolution, regular and irregular structures, terrain models

10 Bifurcated routing in computer networks

Wai Sum Lai
July 1985 **ACM SIGCOMM Computer Communication Review**, Volume 15 Issue 3
Publisher: ACM Press

Full text available: pdf(1.25 MB) Additional Information: full citation, abstract, references

This paper presents a characterization and a survey of multiple path routing in computer networks. It also develops a routing protocol that achieves load sharing and combines the strengths of both virtual circuit and datagram networks.

11 Peer-to-peer systems: Adaptive bitstream switching of pre-encoded PFGS video

Osama Lotfallah, Martin Reisslein, Sethuraman Panchanathan
November 2005 **Proceedings of the ACM workshop on Advances in peer-to-peer multimedia streaming P2PMMS'05**
Publisher: ACM Press
Full text available: pdf(357.39 KB) Additional Information: full citation, abstract, references, index terms

With Progressive Fine Granularity Scalability (PFGS) video coding, one given encoding (with a prescribed bit rate) can flexibly be transmitted at any lower bit rate. However, the transmitted video is only efficiently encoded when the transmission bit rate is in the vicinity of the encoding bit rate; for transmission bit rates far from the encoding bit rate up on the order of 4 dB in video quality are lost. In this paper we develop and evaluate a suite of policies for accounting for this coding e ...

Keywords: PFGS, content-dependent coding, content-dependent packet drop, motion activity level, quality-dependent multicast

12 Demonstration session 2: Identifying audio clips with *RARE*
Chris J. C. Burges, John C. Platt, Jonathan Goldstein
November 2003 **Proceedings of the eleventh ACM international conference on Multimedia MULTIMEDIA '03**
Publisher: ACM Press
Full text available: pdf(159.91 KB) Additional Information: full citation, abstract, references, citings, index terms

In this paper, we describe RARE (Robust Audio Recognition Engine): a system for identifying audio streams and files. RARE can be used in a variety of applications: from enhancing the consumer listening experience to cleaning large audio databases. RARE was designed with two key qualities in mind: robustness to distortion of the audio, and lookup speed. RARE identifies audio clips in a stream against a database of 1/4 million songs in real time using approximately 10% CPU on an 850 MHz P3, and wi ...

Keywords: audio fingerprinting, fast indexing, robust lookup

13 DB-2 (databases): data streams: Interval query indexing for efficient stream processing
Kun-Lung Wu, Shyh-Kwei Chen, Philip S. Yu
November 2004 **Proceedings of the thirteenth ACM international conference on Information and knowledge management CIKM '04**
Publisher: ACM Press
Full text available: pdf(263.69 KB) Additional Information: full citation, abstract, references, citings, index terms

A large number of continual range queries can be issued against a data stream. Usually, a main memory-based query index with a small storage cost and a fast search time is needed, especially if the stream is rapid. In this paper, we present a CEI-based query index that meets both criteria for efficient processing of continual interval queries in a streaming environment. This new query index is centered around a set of predefined virtual containment-encoded intervals, or CEIs. T ...

Keywords: continual queries, data streams, interval indexing, query indexing, query monitoring

14 Spatial queries in wireless broadcast systems
Baihua Zheng, Wang-Chien Lee, Dik Lun Lee

November 2004 **Wireless Networks**, Volume 10 Issue 6
Publisher: Kluwer Academic Publishers
Full text available: pdf(402.37 KB) Additional Information: full citation, abstract, references, citings, index terms

Owing to the advent of wireless networking and personal digital devices, information systems in the era of mobile computing are expected to be able to handle a tremendous amount of traffic and service requests from the users. Wireless data broadcast, thanks to its high scalability, is particularly suitable for meeting such a challenge. Indexing techniques have been developed for wireless data broadcast systems in order to conserve the scarce power resources in mobile clients. However, most of ...

Keywords: index structure, location-dependent spatial queries, pervasive computing, wireless broadcast

15 Research sessions: continuous queries and streams: Continuously adaptive continuous queries over streams
Samuel Madden, Mehul Shah, Joseph M. Hellerstein, Vijayshankar Raman
June 2002 **Proceedings of the 2002 ACM SIGMOD international conference on Management of data SIGMOD '02**
Publisher: ACM Press
Full text available: pdf(1.59 MB) Additional Information: full citation, abstract, references, citings, index terms

We present a continuously adaptive, continuous query (CACQ) implementation based on the eddy query processing framework. We show that our design provides significant performance benefits over existing approaches to evaluating continuous queries, not only because of its adaptivity, but also because of the aggressive cross-query sharing of work and space that it enables. By breaking the abstraction of shared relational algebra expressions, our Telegraph CACQ implementation is able to share physica ...

16 Real-time shading
Marc Olano, Kurt Akeley, John C. Hart, Wolfgang Heidrich, Michael McCool, Jason L. Mitchell, Randi Rost
August 2004 **ACM SIGGRAPH 2004 Course Notes SIGGRAPH '04**
Publisher: ACM Press
Full text available: pdf(7.39 MB) Additional Information: full citation, abstract

Real-time procedural shading was once seen as a distant dream. When the first version of this course was offered four years ago, real-time shading was possible, but only with one-of-a-kind hardware or by combining the effects of tens to hundreds of rendering passes. Today, almost every new computer comes with graphics hardware capable of interactively executing shaders of thousands to tens of thousands of instructions. This course has been redesigned to address today's real-time shading capabili ...

17 Smalltalk-80: the language and its implementation
Adele Goldberg, David Robson
January 1983 Book
Publisher: Addison-Wesley Longman Publishing Co., Inc.
Full text available: pdf(33.56 MB) Additional Information: full citation, abstract, cited by, index terms, review

From the Preface (See Front Matter for full Preface)

Advances in the design and production of computer hardware have brought many more people into direct contact with computers. Similar advances in the design and production of computer software are required in order that this increased contact be as rewarding as

possible. The Smalltalk-80 system is a result of a decade of research into creating computer software that is appropriate for producing highly functional and interactive ...

18 The multics system: an examination of its structure

Elliott I. Organick
January 1972 Book
Publisher: MIT Press
Additional Information: full citation, abstract, references, cited by, index terms

This volume provides an overview of the Multics system developed at M.I.T.--a time-shared, general purpose utility like system with third-generation software. The advantage that this new system has over its predecessors lies in its expanded capacity to manipulate and file information on several levels and to police and control access to data in its various files. On the invitation of M.I.T.'s Project MAC, Elliott Organick developed over a period of years an explanation of the workings, concep ...

19 Query evaluation techniques for large databases

Goetz Graefe
June 1993 **ACM Computing Surveys (CSUR)**, Volume 25 Issue 2
Publisher: ACM Press
Full text available: pdf(9.37 MB) Additional Information: full citation, abstract, references, citings, index terms, review

Database management systems will continue to manage large data volumes. Thus, efficient algorithms for accessing and manipulating large sets and sequences will be required to provide acceptable performance. The advent of object-oriented and extensible database systems will not solve this problem. On the contrary, modern data models exacerbate the problem: In order to manipulate large sets of complex objects as efficiently as today's database systems manipulate simple records, query-processi ...

Keywords: complex query evaluation plans, dynamic query evaluation plans, extensible database systems, iterators, object-oriented database systems, operator model of parallelization, parallel algorithms, relational database systems, set-matching algorithms, sort-hash duality

20 The elements of nature: interactive and realistic techniques

Oliver Deusen, David S. Ebert, Ron Fedkiw, F. Kenton Musgrave, Przemyslaw Prusinkiewicz, Doug Roble, Jos Stam, Jerry Tessendorf
August 2004 **ACM SIGGRAPH 2004 Course Notes SIGGRAPH '04**
Publisher: ACM Press
Full text available: pdf(17.65 MB) Additional Information: full citation, abstract

This updated course on simulating natural phenomena will cover the latest research and production techniques for simulating most of the elements of nature. The presenters will provide movie production, interactive simulation, and research perspectives on the difficult task of photorealistic modeling, rendering, and animation of natural phenomena. The course offers a nice balance of the latest interactive graphics hardware-based simulation techniques and the latest physics-based simulation techni ...

Terms used: **generate bit stream index tree** Found **18,663** of **210,468**

Sort results by: relevance Save results to a Binder

Display results: expanded form Search Tips

Open results in a new window

Try an Advanced Search

Try this search in The ACM Guide

Results 1 - 20 of 200 Result page: **1** 2 3 4 5 6 7 8 9 10 next

Best 200 shown Relevance scale

1 The hB-tree: a multiattribute indexing method with good guaranteed performance

David B. Lomet, Betty Salzberg

December 1990 **ACM Transactions on Database Systems (TODS)**, Volume 15 Issue 4

Publisher: ACM Press

Full text available: pdf(2.58 MB) Additional Information: full citation, abstract, references, citings, index terms, review

A new multiattribute index structure called the hB-tree is introduced. It is derived from the K-D-B-tree of Robinson [15] but has additional desirable properties. The hB-tree internode search and growth processes are precisely analogous to the corresponding processes in B-trees [1]. The intranode processes are unique. A k-d tree is used as the structure within nodes for very efficient searching. Node splitting requires that this k-d tree be split. This produces nodes which no longer represe ...

2 RE-tree: an efficient index structure for regular expressions

Chee-Yong Chan, Minos Garofalakis, Rajeev Rastogi

August 2003 **The VLDB Journal — The International Journal on Very Large Data Bases**, Volume 12 Issue 2

Publisher: Springer-Verlag New York, Inc.

Full text available: pdf(346.00 KB) Additional Information: full citation, abstract, index terms

Abstract.Due to their expressive power, regular expressions (REs) are quickly becoming an integral part of language specifications for several important application scenarios. Many of these applications have to manage huge databases of RE specifications and need to provide an effective matching mechanism that, given an input string, quickly identifies the REs in the database that match it. In this paper, we propose the RE-tree, a novel index structure for large databases of RE specifications. Gi ...

Keywords: Index structure, Regular expressions, Sampling-based approximations, Size measures

3 XML indexing and compression: ViST: a dynamic index method for querying XML data by tree structures

Haixun Wang, Sanghyun Park, Wei Fan, Philip S. Yu

June 2003 **Proceedings of the 2003 ACM SIGMOD international conference on Management of data SIGMOD '03**

Publisher: ACM Press

Full text available: pdf(244.47 KB) Additional Information: full citation, abstract, references, citings, index terms

With the growing importance of XML in data exchange, much research has been done in providing flexible query facilities to extract data from structured XML documents. In this paper, we propose ViST, a novel index structure for searching XML documents. By representing both XML documents and XML queries in structure-encoded sequences, we show that querying XML data is equivalent to finding subsequence matches. Unlike index methods that disassemble a query into multiple sub-queries, and then *joi* ...

4 Database principles: A mapping mechanism to support bitmap index and other auxiliary structures on tables stored as primary B$^+$-trees

Eugene Inseok Chong, Jagannathan Srinivasan, Souripriya Das, Chuck Freiwald, Aravind Yalamanchi, Mahesh Jagannath, Anh-Tuan Tran, Ramkumar Krishnan, Richard Jiang
June 2003 **ACM SIGMOD Record**, Volume 32 Issue 2
Publisher: ACM Press
Full text available: pdf(198.55 KB) Additional Information: full citation, abstract, references

Any auxiliary structure, such as a bitmap or a B$^+$-tree index, that refers to rows of a table stored as a primary B$^+$-tree (e.g., *tables with clustered index* in Microsoft SQL Server, or *index-organized tables* in Oracle) by their physical addresses would require updates due to inherent volatility of those addresses. To address this problem, we propose a mapping mechanism that 1) introduces a single mapping table, with each row holding one key value from the prima ...

5 Concurrency and recovery for index trees

David Lomet, Betty Salzberg
August 1997 **The VLDB Journal — The International Journal on Very Large Data Bases**, Volume 6 Issue 3
Publisher: Springer-Verlag New York, Inc.
Full text available: pdf(168.36 KB) Additional Information: full citation, abstract, citings, index terms

Although many suggestions have been made for concurrency in B $^+$-trees, few of these have considered recovery as well. We describe an approach which provides high concurrency while preserving well-formed trees across system crashes. Our approach works for a class of index trees that is a generalization of the B $^{\rm link}$-tree. This class includes some multi-attribute indexes and temporal indexes. Structural changes in an index tree are decomposed into a sequence of atomic actions, each one ...

Keywords: Access methods, B-trees, Concurrency, Indexing, Recovery

6 The hB $^\Pi$-tree: a multi-attribute index supporting concurrency, recovery and node consolidation

Georgios Evangelidis, David Lomet, Betty Salzberg
February 1997 **The VLDB Journal — The International Journal on Very Large Data Bases**, Volume 6 Issue 1
Publisher: Springer-Verlag New York, Inc.
Full text available: pdf(314.15 KB) Additional Information: full citation, abstract, citings, index terms

We propose a new multi-attribute index. Our approach combines the hB-tree, a multi-attribute index, and the Π-tree, an abstract index which offers efficient concurrency and recovery methods. We call the resulting method the hB $^\Pi$-tree. We describe several versions of the hB $^\Pi$-tree, each using a different node-splitting and index-term-posting algorithm. We also describe a new node deletion algorithm. We have implemented all the versions of the hB $^\Pi$-tree. Our performance results ...

Keywords: Concurrency, Multi-attribute index, Node consolidation, Recovery

7 Dynamic vp-tree indexing for *n*-nearest neighbor search given pair-wise distances

Ada Wai-chee Fu, Polly Mei-shuen Chan, Yin-Ling Cheung, Yiu Sang Moon

July 2000 **The VLDB Journal — The International Journal on Very Large Data Bases**, Volume 9 Issue 2

Publisher: Springer-Verlag New York, Inc.

Full text available: pdf(232.09 KB) Additional Information: full citation, abstract, citings, index terms

For some multimedia applications, it has been found that domain objects cannot be represented as feature vectors in a multidimensional space. Instead, pair-wise distances between data objects are the only input. To support content-based retrieval, one approach maps each object to a *k*-dimensional (*k*-d) point and tries to preserve the distances among the points. Then, existing spatial access index methods such as the R-trees and KD-trees can support fast searching on the resulting

***Keywords**: Content-based retrieval, Indexing, Nearest neighbor search, Pair-wise distances, Updating*

8 Industry session 3: database performance and interface: A mapping mechanism to support bitmap index and other auxiliary structures on tables stored as primary B^{+}-trees

Eugene Inseok Chong, Jagannathan Srinivasan, Souripriya Das, Chuck Freiwald, Aravind Yalamanchi, Mahesh Jagannath, Anh-Tuan Tran, Ramkumar Krishnan, Richard Jiang

November 2002 **Proceedings of the eleventh international conference on Information and knowledge management CIKM '02**

Publisher: ACM Press

Full text available: pdf(63.19 KB) Additional Information: full citation, abstract, references, index terms

Any auxiliary structure, such as a bitmap or a B^{+}-tree index, that refers to rows of a table stored as a primary B^{+}-tree (e.g., *tables with clustered index* in Microsoft SQL Server, or *index-organized tables* in Oracle) by their physical addresses would require updates due to inherent volatility of those addresses. To address this problem, we propose a mapping mechanism that 1) introduces a single *mapping table*, with each row holding one key value from th ...

Keywords: bitmap indexes, mapping mechanism, primary B^{+}-trees

9 Industrial sessions: commercial implementation techniques: Quadtree and R-tree indexes in oracle spatial: a comparison using GIS data

Ravi Kanth V Kothuri, Siva Ravada, Daniel Abugov

June 2002 **Proceedings of the 2002 ACM SIGMOD international conference on Management of data SIGMOD '02**

Publisher: ACM Press

Full text available: pdf(1.03 MB) Additional Information: full citation, abstract, references, citings, index terms

Spatial indexing has been one of the active focus areas in recent database research. Several variants of Quadtree and R-tree indexes have been proposed in database literature. In this paper, we first describe briefly our implementation of Quadtree and R-tree index structures and related optimizations in Oracle Spatial. We then examine the relative merits of two structures as implemented in Oracle Spatial and compare their performance for different types of queries and other operations. Finally, ...

10 Optimizing multidimensional index trees for main memory access

Kihong Kim, Sang K. Cha, Keunjoo Kwon
May 2001 **ACM SIGMOD Record , Proceedings of the 2001 ACM SIGMOD international conference on Management of data SIGMOD '01**, Volume 30 Issue 2
Publisher: ACM Press

Full text available: pdf(243.75 KB) Additional Information: full citation, abstract, references, citings, index terms

Recent studies have shown that cache-conscious indexes such as the CSB+-tree outperform conventional main memory indexes such as the T-tree. The key idea of these cache-conscious indexes is to eliminate most of child pointers from a node to increase the fanout of the tree. When the node size is chosen in the order of the cache block size, this pointer elimination effectively reduces the tree height, and thus improves the cache behavior of the index. However, the pointer elimination cannot be ...

11 Isosurface extraction in time-varying fields using a temporal hierarchical index tree
Han-Wei Shen
October 1998 **Proceedings of the conference on Visualization '98 VIS '98**
Publisher: IEEE Computer Society Press
Full text available: pdf(1.24 MB) Publisher Site Additional Information: full citation, references, citings, index terms

Keywords: isosurface extraction, marching cubes, scalar field visualization, span space, time-varying fields, volume visualization

12 The PML-tree: an efficient parallel spatial index structure for spatial databases
Kap S. Bang, Huizhu Lu
February 1996 **Proceedings of the 1996 ACM 24th annual conference on Computer science CSC '96**
Publisher: ACM Press
Full text available: pdf(903.84 KB) Additional Information: full citation, references, index terms

13 H-trees: a dynamic associative search index for OODB
Chee Chin Low, Beng Chin Ooi, Hongjun Lu
June 1992 **ACM SIGMOD Record , Proceedings of the 1992 ACM SIGMOD international conference on Management of data SIGMOD '92**, Volume 21 Issue 2
Publisher: ACM Press

Full text available: pdf(1.06 MB) Additional Information: full citation, abstract, references, citings, index terms

The support of the superclass-subclass concept in object-oriented databases (OODB) makes an instance of a subclass also an instance of its superclass. As a result, the access scope of a query against a class in general includes the access scope of all its subclasses, unless specified otherwise. To support the superclass-subclass relationship efficiently, the index must achieve two objectives. First, the index must support efficient retrieval of instances from a single class. Second, it must ...

14 Join processing and indexing: A combination of trie-trees and inverted files for the indexing of set-valued attributes
Manolis Terrovitis, Spyros Passas, Panos Vassiliadis, Timos Sellis
November 2006 **Proceedings of the 15th ACM international conference on Information and knowledge management CIKM '06**
Publisher: ACM Press

Full text available: pdf(516.40 KB) Additional Information: full citation, abstract, references, index terms

Set-valued attributes frequently occur in contexts like market-basked analysis and stock market trends. Late research literature has mainly focused on set containment joins and data mining without considering simple queries on set valued attributes. In this paper we address superset, subset and equality queries and we propose a novel indexing scheme for answering them on set-valued attributes. The proposed index superimposes a trie-tree on top of an inverted file that indexes a relation with set ...

Keywords: HTI, containment queries, inverted files, tries

15 Graphs and trees: Heuristic containment check of partial tree-pattern queries in the presence of index graphs

Dimitri Theodoratos, Stefanos Souldatos, Theodore Dalamagas, Pawel Placek, Timos Sellis
November 2006 **Proceedings of the 15th ACM international conference on Information and knowledge management CIKM '06**
Publisher: ACM Press
Full text available: pdf(281.90 KB) Additional Information: full citation, abstract, references, index terms

The wide adoption of XML has increased the interest of the database community on tree-structured data management techniques. Querying capabilities are provided through tree-pattern queries. The need for querying tree-structured data sources when their structure is not fully known, and the need to integrate multiple data sources with different tree structures have driven, recently, the suggestion of query languages that relax the complete specification of a tree pattern. In this paper, we use a q ...

Keywords: partial tree-pattern query, query containment, tree-structured data

16 iDistance: An adaptive B+-tree based indexing method for nearest neighbor search

H. V. Jagadish, Beng Chin Ooi, Kian-Lee Tan, Cui Yu, Rui Zhang
June 2005 **ACM Transactions on Database Systems (TODS)**, Volume 30 Issue 2
Publisher: ACM Press
Full text available: pdf(1.16 MB) Additional Information: full citation, abstract, references, citings, index terms

In this article, we present an efficient B+-tree based indexing method, called iDistance, for K-nearest neighbor (KNN) search in a high-dimensional metric space. iDistance partitions the data based on a space- or data-partitioning strategy, and selects a reference point for each partition. The data points in each partition are transformed into a single dimensional value based on their similarity with respect to the reference point. This allows the points to be indexed using a B
~~**Keywords**: Indexing, KNN, nearest neighbor queries~~

17 Research papers: storage, indexing, and system architecture: Online B-tree merging

Xiaowei Sun, Rui Wang, Betty Salzberg, Chendong Zou
June 2005 **Proceedings of the 2005 ACM SIGMOD international conference on Management of data SIGMOD '05**
Publisher: ACM Press
Full text available: pdf(394.41 KB) Additional Information: full citation, abstract, references

Many scenarios involve merging of two B-tree indexes, both covering the same key range. Increasing demand for continuous availability and high performance requires that such merging be done online, with minimal interference to normal user transactions. In this paper we present an online B-tree merging method, in which the merging of leaf pages in

two B-trees are piggybacked lazily with normal user transactions, thus making the merging I/O efficient and allowing user transactions to access only o ...

18 BRep-Index: a multidimensional space partitioning tree

George Vaněček

May 1991 **Proceedings of the first ACM symposium on Solid modeling foundations and CAD/CAM applications SMA '91**

Publisher: ACM Press

Full text available: pdf(911.29 KB) Additional Information: full citation, references, citings, index terms

19 Physical storage structures: B+ trees and indexed sequential files: a performance comparison

D. S. Batory

April 1981 **Proceedings of the 1981 ACM SIGMOD international conference on Management of data SIGMOD '81**

Publisher: ACM Press

Full text available: pdf(858.82 KB) Additional Information: full citation, abstract, references, citings

An analytic method for comparing the performance of B+ trees and indexed sequential files is proposed. Preliminary results indicate that indexed sequential files may be more efficient than B+ trees in certain applications.

20 The SR-tree: an index structure for high-dimensional nearest neighbor queries

Norio Katayama, Shin'ichi Satoh

June 1997 **ACM SIGMOD Record , Proceedings of the 1997 ACM SIGMOD international conference on Management of data SIGMOD '97**, Volume 26 Issue 2

Publisher: ACM Press

Full text available: pdf(1.41 MB) Additional Information: full citation, abstract, references, citings, index terms

Recently, similarity queries on feature vectors have been widely used to perform content-based retrieval of images. To apply this technique to large databases, it is required to develop multidimensional index structures supporting nearest neighbor queries efficiently. The SS-tree had been proposed for this purpose and is known to outperform other index structures such as the R*-tree and the K-D-B-tree. One of its most important features is that it employs bounding spheres rather than boundi ...

Results 1 - 20 of 200 Result page: **1** 2 3 4 5 6 7 8 9 10 next